

Learning to Control a Brain–Machine Interface for Reaching and Grasping by Primates

Jose M. Carmena^{1,4}, Mikhail A. Lebedev^{1,4}, Roy E. Crist¹, Joseph E. O’Doherty², David M. Santucci¹, Dragan F. Dimitrov^{1,3}, Parag G. Patil^{1,3}, Craig S. Henriquez^{2,4}, Miguel A. L. Nicolelis^{1,2,4,5*}

1 Department of Neurobiology, Duke University, Durham, North Carolina, United States of America, **2** Department of Biomedical Engineering, Duke University, Durham, North Carolina, United States of America, **3** Division of Neurosurgery, Duke University, Durham, North Carolina, United States of America, **4** Center for Neuroengineering, Duke University, Durham, North Carolina, United States of America, **5** Department of Psychological and Brain Sciences, Duke University, Durham, North Carolina, United States of America

Reaching and grasping in primates depend on the coordination of neural activity in large frontoparietal ensembles. Here we demonstrate that primates can learn to reach and grasp virtual objects by controlling a robot arm through a closed-loop brain–machine interface (BMIc) that uses multiple mathematical models to extract several motor parameters (i.e., hand position, velocity, gripping force, and the EMGs of multiple arm muscles) from the electrical activity of frontoparietal neuronal ensembles. As single neurons typically contribute to the encoding of several motor parameters, we observed that high BMIc accuracy required recording from large neuronal ensembles. Continuous BMIc operation by monkeys led to significant improvements in both model predictions and behavioral performance. Using visual feedback, monkeys succeeded in producing robot reach-and-grasp movements even when their arms did not move. Learning to operate the BMIc was paralleled by functional reorganization in multiple cortical areas, suggesting that the dynamic properties of the BMIc were incorporated into motor and sensory cortical representations.

Introduction

Traumatic lesions of the central nervous system as well as neurodegenerative disorders continue to inflict devastating, and so far irreparable, motor deficits in large numbers of patients. Every year, spinal cord injuries alone are responsible for the occurrence of about 11,000 new cases of permanent paralysis in the United States (Nobunaga et al. 1999). These cases add up to an already sizeable population of patients, estimated at 200,000 in the United States (Nobunaga et al. 1999), who have to cope with partial (as in the case of paraplegics) or almost total (i.e., quadriplegia) body paralysis.

Until very recently, the main thrust of basic research on restoration of motor functions after spinal cord injuries focused on reconstructing the connectivity and functionality of damaged nerve fibers (Ramon-Cueto et al. 1998; Uchida et al. 2000; Bomze et al. 2001; Bunge 2001; Schwab 2002). While this repair strategy has produced encouraging results, such as limited restoration of limb mobility in animals, the goal of restoring complex motor behaviors, such as reaching and grasping, remains a major challenge.

Two decades ago, an alternative method for restoring motor behaviors in severely paralyzed patients was proposed (Schmidt 1980). This approach contends that direct interfaces between spared cortical or subcortical motor centers and artificial actuators could be employed to “bypass” spinal cord injuries so that paralyzed patients could enact their voluntary motor intentions. Initial experimental support for a cortically driven bypass came from the studies conducted by Fetz and collaborators (Fetz 1969; Fetz and Finocchio 1971, 1975; Fetz and Baker 1973), who demonstrated that macaque monkeys could learn to selectively adjust the firing rate of individual cortical neurons to attain a particular level of cell activity if provided with sensory feedback that signaled the level of neuronal firing.

Recent studies in rodents (Chapin et al. 1999; Talwar et al. 2002), primates (Wessberg et al. 2000; Serruya et al. 2002;

Taylor et al. 2002), and human subjects (Birbaumer 1999) have rekindled interest in using brain–machine interfaces (BMIs) as a potential alternative for spinal cord rehabilitation. These experiments have demonstrated that animals can learn to utilize their brain activity to control the displacements of computer cursors (Serruya et al. 2002; Taylor et al. 2002) or one-dimensional (1D) to three-dimensional (3D) movements of simple and elaborate robot arms (Chapin et al. 1999; Wessberg et al. 2000).

Despite these initial results, several fundamental issues regarding the operation of BMIs, ranging from basic electrophysiological issues to multiple engineering bottlenecks, remain a matter of considerable debate (Nicolelis 2001, 2003; Donoghue 2002). For example, although most agree that a BMI designed to reproduce arm/hand movements will require long-term and stable recordings from cortical (or subcortical) neurons through chronically implanted electrode arrays (Nicolelis 2001, 2003; Donoghue 2002), there is considerable disagreement on what type of brain signal (single unit, multiunit, or field potentials [Pesaran et al. 2002])

Received August 14, 2003; Accepted September 3, 2003; Published October 13, 2003

DOI: 10.1371/journal.pbio.0000042

Copyright: © 2003 Carmena et al. This is an open-access article distributed under the terms of the Public Library of Science Open-Access License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abbreviations: BMI, brain–machine interface; BMIc, closed-loop brain–machine interface; 1D, one dimensional; 3D, three dimensional; DOF, degree of freedom; DTC, directional tuning curve; DTD, directional tuning depth; DTE, directional tuning of neural ensemble; EMG, electromyogram; GF, gripping force; HP, hand position; HV, hand velocity; M1, primary motor cortex; MIP, medial intraparietal area; ND, neuron dropping; PETH, perievent time histogram; PMd, dorsal premotor cortex; PP, posterior parietal cortex; S1, primary somatosensory cortex; SMA, supplementary motor area

Academic Editor: Idan Segev, Hebrew University

*To whom correspondence should be addressed. E-mail: nicoleli@neuro.duke.edu



would be the optimal input for a such a device. Research groups that propose the use of single-unit activity also diverge on the assessments of whether small (eight to thirty) (Serruya et al. 2002; Taylor et al. 2002) or substantially larger (hundreds to thousands) numbers of single units may be necessary to operate a BMI efficiently for many years (Wessberg et al. 2000; Nicolelis 2001).

The issues of what signal to use and how much neuronal tissue to sample are linked to the question of what type of motor commands may be extracted from brain activity. Up to now, animal studies have demonstrated the capability of extracting a single motor parameter (e.g., hand trajectory, position, direction of movement) from brain activity in order to operate a BMI (Wessberg et al. 2000; Taylor et al. 2002). While it is well known that cortical neuronal activity can encode a variety of motor parameters (Fetz 1992; Messier and Kalaska 2000; Johnson 2001), it is not clear which cortical areas would provide the best input for a BMI designed to restore multiple features of upper-limb function. A couple of laboratories have focused on the primary motor cortex (M1) (Serruya et al. 2002; Taylor et al. 2002), while another group has selected the parietal cortex as the main input to a BMI (Pesaran et al. 2002). Our previous studies have suggested that, because of the distributed nature of motor planning in the brain, neuronal samples from multiple frontal and parietal cortical areas ought to be employed to operate such devices (Wessberg et al. 2000; Nicolelis 2001, 2003; Donoghue 2002). Another important issue that has received little attention is how the interposition of an artificial actuator (such as a robot arm) in the control loop impacts the BMI and the subject's performance. Two previous studies have reported that macaque monkeys learn to operate a closed-loop BMI (BMiC) using visual feedback (Serruya et al. 2002; Taylor et al. 2002), but the animals in these studies did not control a real mechanical actuator.

Finally, more data are needed to evaluate the extent, relevance, and behavioral meaning of cortical reorganization that can be triggered by operation of a BMiC. A possibility of such reorganization is supported by results in plasticity in M1 neurons in a force-field adaptation task (Li et al. 2001) and by an initial report of changes in directional selectivity in a small sample of M1 neurons during BMiC operation (Taylor et al. 2002).

In this paper, we present the results from a series of long-term studies in macaque monkeys to address several of the fundamental issues that currently shape the debate on BMIs. In particular, we demonstrate for what we believe is the first time the ability of the same ensemble of cells in closed-loop mode to control two distinct movements of a robotic arm: reaching and grasping. In addition, we demonstrate how the monkeys learn to control a real robotic actuator using a BMiC. We also report on how they overcome the robot dynamics and return to the same level of performance without modification of the task. Finally, we quantitatively compare the contribution of neural populations in multiple cortical areas needed to create this control and analyze changes in these contributions during learning.

Results

Using the experimental apparatus illustrated in Figure 1A, monkeys were trained in three different tasks: a reaching task

(task 1; Figure 1B), a hand-gripping task (task 2; Figure 1B), and a reach-and-grasp task (task 3; Figure 1B).

We used multiple linear models, similar to those described in our previous studies (Wessberg et al. 2000), to simultaneously extract a variety of motor parameters (i.e., hand position [HP_x, HP_y, HP_z], velocity [HV_x, HV_y, HV_z], and gripping force [GF]) and multiple muscle electromyograms (EMGs) from the activity of cortical neural ensembles. Although all these parameters were extracted in real time on each session, only some of them were used to control the BMiC, depending on each of the three tasks the monkeys had to solve in a given day. In each recording session, an initial 30-min period was used for training of these models. During this period, monkeys used a hand-held pole either to move a cursor on the screen or to change the cursor size by application of gripping force to the pole. This period is referred to as “pole control” mode. As the models converged to an optimal performance, their coefficients were fixed and the control of the cursor position (task 1 and 3) and/or size (task 2 and 3) was obtained directly from the output of the linear models. This period is referred to as “brain control” mode. During brain control mode, animals initially produced arm movements, but they soon realized that these were not necessary and ceased to produce them for periods of time. To systematically study this phenomenon, we removed the pole after the monkey ceased to produce arm movements in a session. In each task, after initial training, a 6 DOF (degree-of-freedom) robot arm equipped with a 1 DOF gripper was included in the BMiC control loop. In all experiments, visual feedback (i.e., cursor position/size) informed the animal about the BMiC's performance. When the robot was used, cursor position indicated to the animal the X and Y coordinates of the robot hand. The cursor size provided feedback of the force measured by the sensors on the robot's gripper. The time delay between the output of the linear model and the response of the robot was in the range of 60–90 ms.

For each task, training continued until the animal reached high levels of performance in brain control mode. During this learning period, both the animal's and the BMiC's performance were assessed using several measures. Chance performance was assessed using Monte Carlo simulations of random walks. Contributions of individual neurons and the overall contribution of different cortical areas to the prediction of multiple motor signals were evaluated. In addition, changes in directional tuning of the neurons that resulted from using the BMiC were quantified.

Behavioral Performance during Long-Term Operation of a BMiC

Figure 1C–1E illustrates procedural motor learning as animals interacted with the BMiC in each of the three tasks. Improvement in behavioral performance with the BMiC was indicated by a significant increase in the percentage of trials completed successfully (Figure 1C–1E, top graphs) and by a reduction in movement time (Figure 1C–1E, bottom graphs). For task 1, both monkeys had some training in pole control of task 1 (data not shown) several weeks before the series of successive daily sessions illustrated in Figure 1C. For both tasks 1 and 2, after a relatively small number of daily training sessions, the monkeys' performance in brain control reached levels similar to those during pole control (Figure 1C and 1D). For tasks 2 and 3, all behavioral data are plotted, given that in

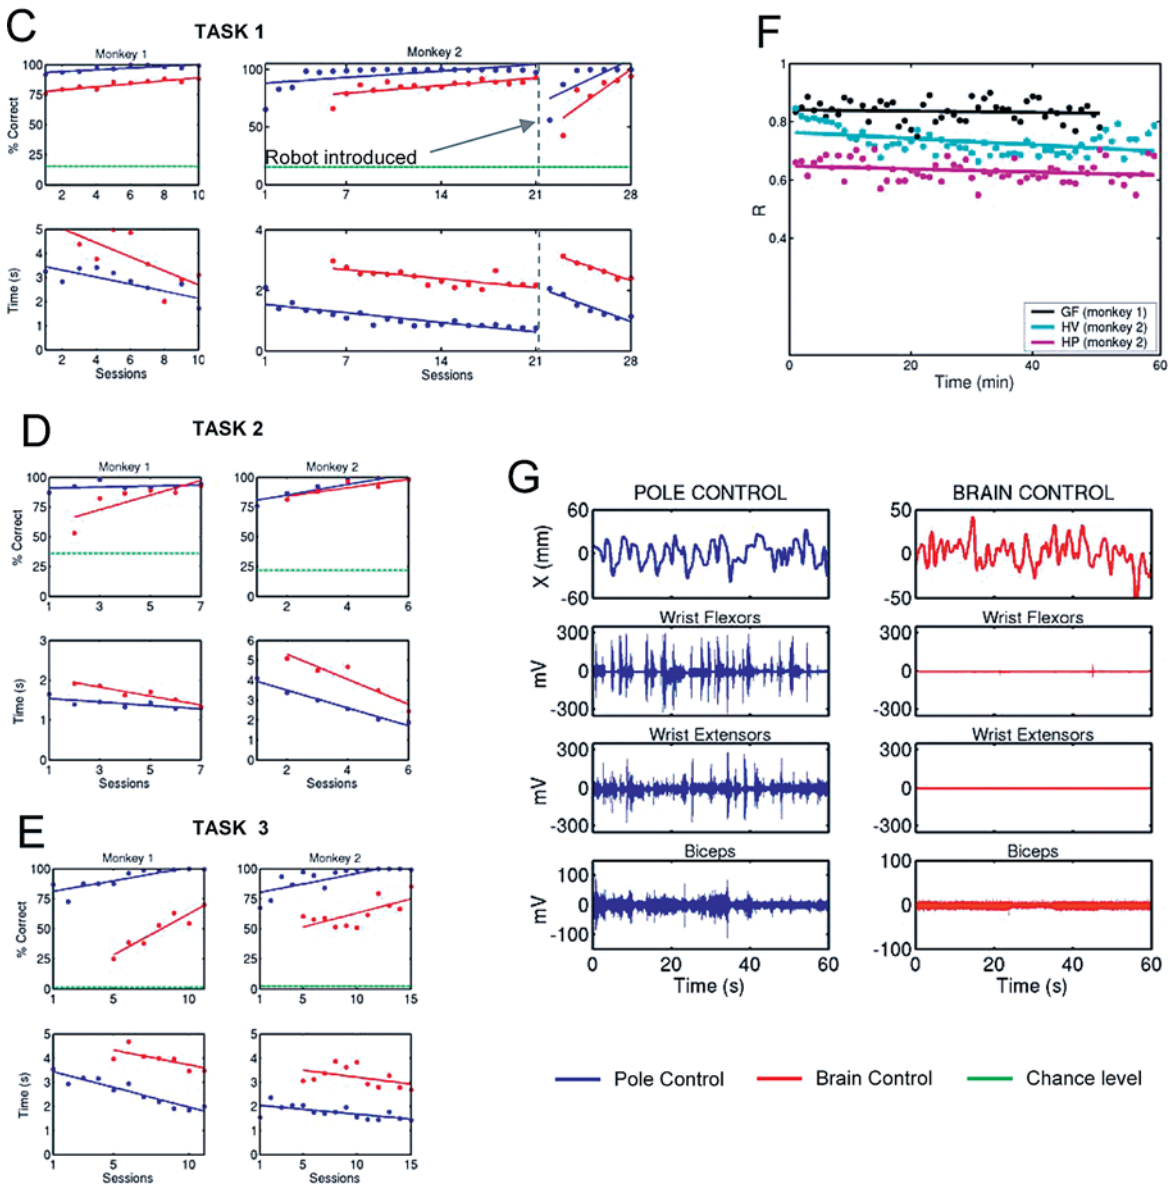
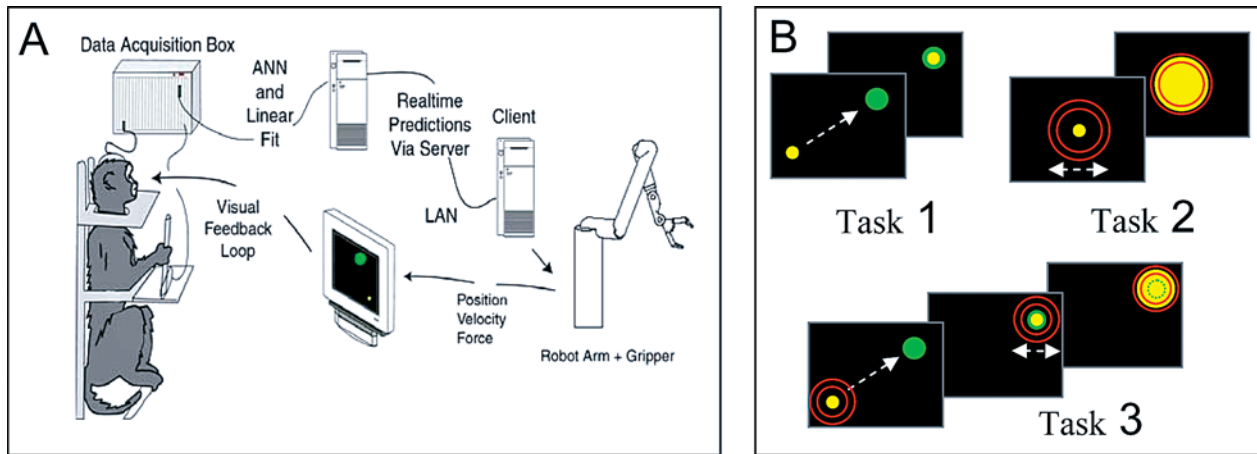


Figure 1. Experimental Setup, Behavioral Tasks, Changes in Performance with Training, EMG Records during Pole and Brain Control, and Stability of Model Predictions

(A) Behavioral setup and control loops, consisting of the data acquisition system, the computer running multiple linear models in real time, the robot arm equipped with a gripper, and the visual display. The pole was equipped with a gripping force transducer. Robot position was translated into cursor position on the screen, and feedback of the gripping force was provided by changing the cursor size.

(B) Schematics of three behavioral tasks. In task 1, the monkey's goal was to move the cursor to a visual target (green) that appeared at random locations on the screen. In task 2, the pole was stationary, and the monkey had to grasp a virtual object by developing a particular gripping force instructed by two red circles displayed on the screen. Task 3 was a combination of tasks 1 and 2. The monkey had to move the cursor to the target and then develop a gripping force necessary to grasp a virtual object.

(C–E) Behavioral performance for two monkeys in tasks 1–3. The percentage of correctly completed trials increased, while the time to conclude a trial decreased with training. This was true for both pole (blue) and brain (red) control. Horizontal (green) lines indicate chance performance obtained from the random walk model. The introduction of the robot arm into the BMIc control loop resulted in a drop in behavioral performance. In approximately seven training sessions, the animal's behavioral performance gradually returned to the initial values. This effect took place during both pole and brain control.

(F) Stability of model predictions of hand velocity during long pole-control sessions (more than 50 min) for two monkeys performing task 1. The first 10 min of performance were used to train the model, and then its coefficients were frozen. Model predictions remained highly accurate for tens of minutes.

(G) Surface EMGs of arm muscles recorded in task 1 for pole control (left) and brain control without arm movements (right). Top plots show the X-coordinate of the cursor; plots below display EMGs of wrist flexors, wrist extensors, and biceps. EMG modulations were absent in brain control.

DOI: 10.1371/journal.pbio.0000042.g001

both cases pole and brain controls were used since the first day of training. Behavioral improvement was also observed in task 3, which combined elements of tasks 1 and 2 (Figure 1E). In all three tasks, the levels of performance attained during brain control mode by far exceeded those predicted by a random walk model (dashed and dotted lines in Figure 1C–1E). Moreover, both animals could operate the BMIc without any overt arm movement and muscle activity, as demonstrated by the lack of EMG activity in several arm muscles (Figure 1G). The ratios of the standard deviation of the muscle activity during pole versus brain control for these muscles were 14.67 (wrist flexors), 9.87 (wrist extensors), and 2.77 (biceps).

A key novel feature of this study was the introduction of the robot equipped with a gripper into the control loop of the BMIc after the animals had learned the task. Figure 1C shows that because the intrinsic dynamics of the robot produced a lag between the pole movement and the cursor movement, the monkeys' performance initially declined. With time, however, the performance rapidly returned to the same levels as seen in previous training sessions (Figure 1C). It is critical to note that the high accuracy in the control of the robot was achieved by using velocity control in the BMIc, which produced smooth predicted trajectories, and by the fine tuning of robot controller parameters. These parameters were fixed across sessions in both monkeys. The controller sent velocity commands to the robot every 60–90 ms. Each of these commands compensated for potential position errors of the robot hand that resulted from previous commands.

In all experiments, the animals continuously received visual feedback of their performance. Unlike previous results in owl monkeys from experiments in which an open-loop BMI was implemented (Wessberg and Nicolelis 2003), after the model parameters were fixed, its predictions did not drift substantially from initial best performance, even during 1-h recordings. As shown in the examples of Figure 1F, prediction of grasping force (mean \pm SEM, $R = 0.84 \pm 5 \times 10^{-3}$) in monkey 1, as well as hand position ($R = 0.63 \pm 3 \times 10^{-3}$) and velocity ($R = 0.73 \pm 5 \times 10^{-3}$) in monkey 2, remained very stable despite some transient fluctuations (slopes for black, magenta, and cyan lines are, respectively, -2.16×10^{-4} , -5.15×10^{-4} , and -1.1×10^{-3}). One possibility is that the presence of continuous visual feedback helped to stabilize model performance.

Which Motor Parameters Can Be Extracted in Real Time?

Throughout learning of all three behavioral tasks, populations of neurons distributed in multiple frontal and parietal cortical areas exhibited task-related modulations of their firing rates. Using multiple linear models running in parallel, several motor signals were extracted from those modulations. To evaluate the performance of the models in extracting different motor parameters, the models were first trained using 15 min of pole control data and then subsequent data were predicted. Figure 2A shows representative 1-min records of such predictions of hand position (HPx, HPy), hand velocity (HVx, HVy), and gripping force (GF). Figure 2B shows the model prediction of EMG activity. In well-trained animals, the linear models accounted for up to 85% of the variance of hand position, 80% of hand velocity, 95% of gripping force, and 61% of multiple EMG activity. These results show that elaborate hand movements, such as the ones required to solve task 3, could be predicted from brain activity using a BMIc with the simultaneous application of multiple linear models.

What Cortical Areas to Sample? How Many Neurons to Record From? What Type of Neuronal Signal to Use?

Several analytical tools were used to address these fundamental questions. By measuring the correlation between neuronal firing and each of the predicted parameters (Figure 2A), we observed that single neurons located in frontal and parietal areas contributed to real-time predictions of all motor parameters analyzed (Figure 2C). Although cortical areas are known to have functional specializations (Wise et al. 1997; Burnod et al. 1999), our sample of M1, dorsal premotor cortex (PMd), supplementary motor area (SMA), posterior parietal cortex (PP), and primary somatosensory cortex (S1) cells provided information, albeit at different levels, for predictions of hand position, velocity, gripping force, and multiple EMGs.

For each motor parameter analyzed, increasing the size of the neuronal population improved the quality of prediction. The effect of sample size on predictions was clearly shown using neuron-dropping (ND) analysis (Figure 2D–2F). ND plots indicate the number of neurons that are required to achieve a particular level of model prediction for each cortical area. Although all cortical areas surveyed contained

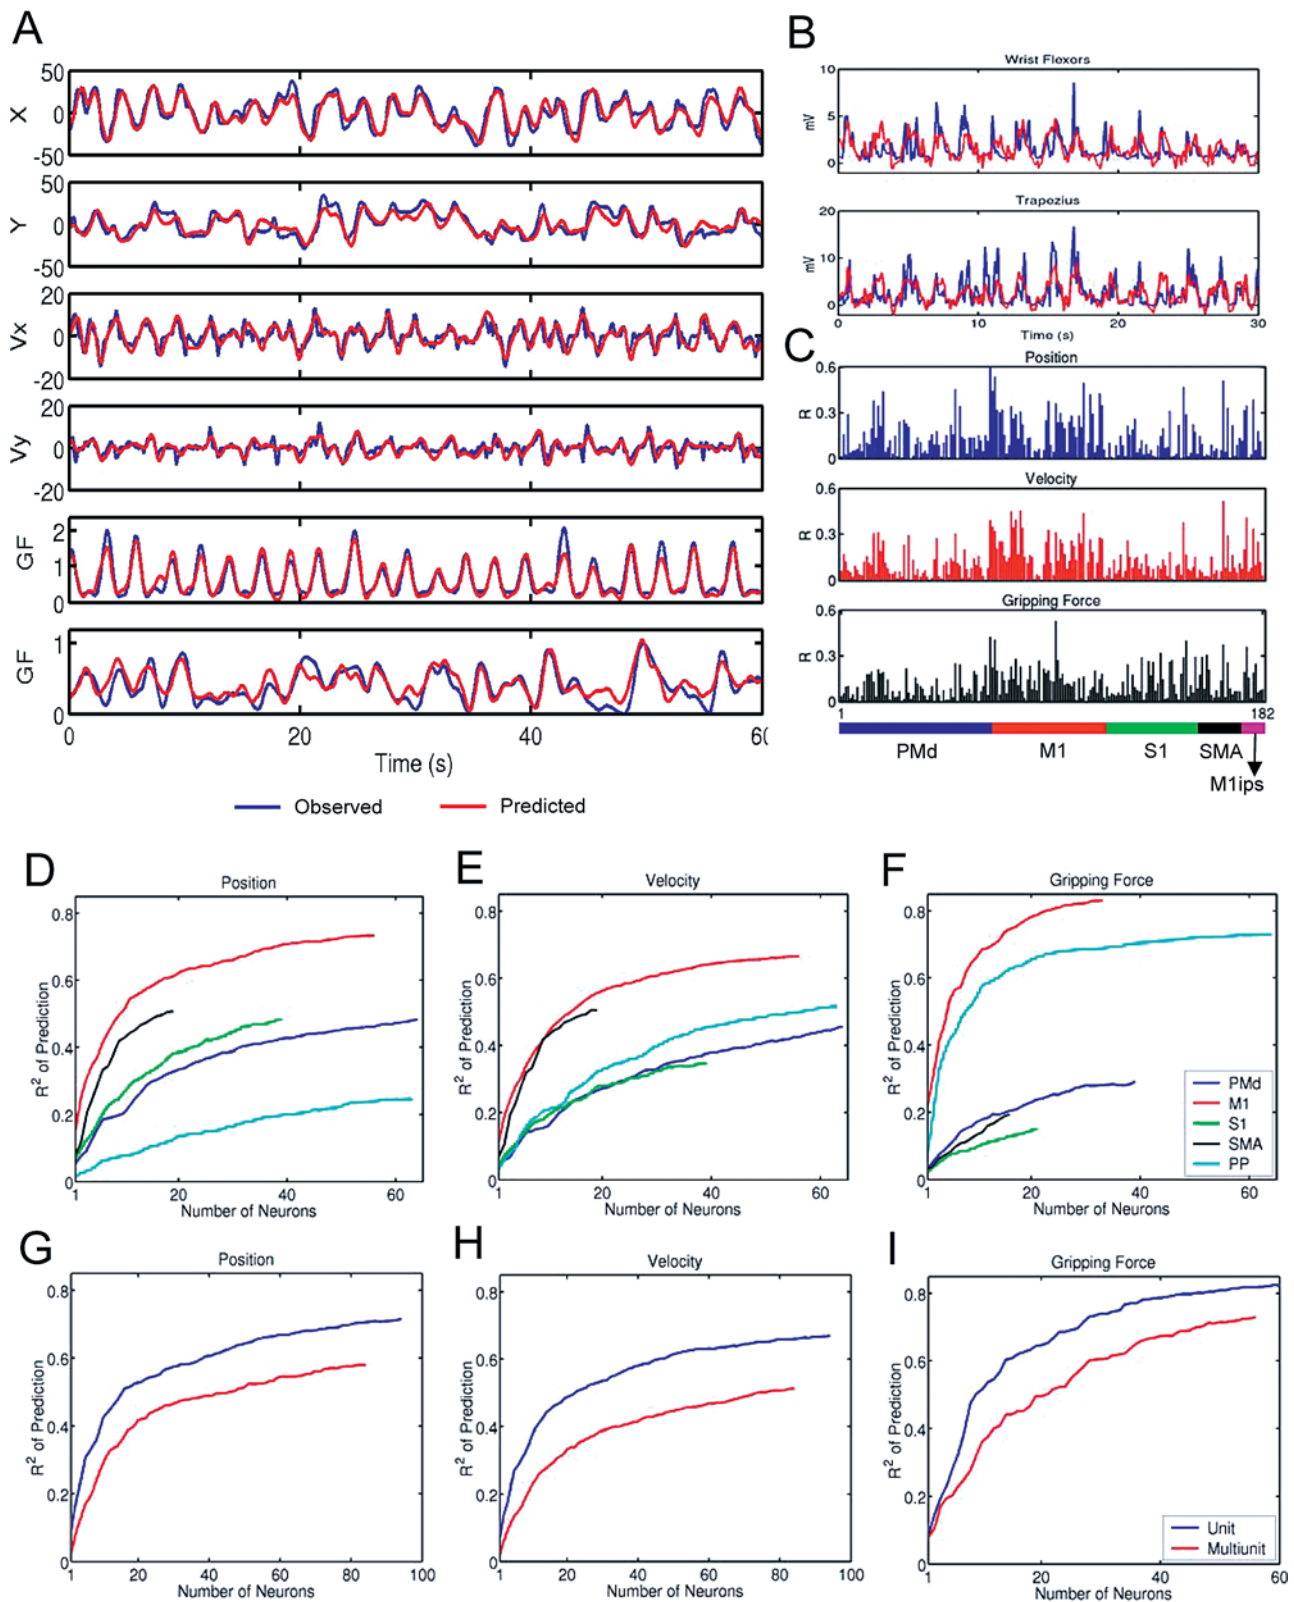


Figure 2. Performance of Linear Models in Predicting Multiple Parameters of Arm Movements, Gripping Force, and EMG from the Activity Frontoparietal Neuronal Ensembles Recorded in Pole Control

(A) Motor parameters (blue) and their prediction using linear models (red). From top to bottom: Hand position (HPx, HPy) and velocity (HVx, HVy) during execution of task 1 and gripping force (GF) during execution of tasks 2 and 1.

(B) EMGs (blue) recorded in task 1 and their prediction (red).

(C) Contribution of neurons from the same ensemble to predictions of hand position (top), velocity (middle), and gripping force (bottom). Contributions were measured as correlation coefficients (R) between the recorded motor parameters and their values predicted by the linear

model. The color bar at the bottom indicates cortical areas where the neurons were located. Each neuron contributed to prediction of multiple parameters of movements, and each area contained information about all parameters. (D–F) Contribution of different cortical areas to model predictions of hand position, velocity (task 1), and gripping force (task 2). For each area, ND curves represent the average prediction accuracy (R^2) as a function of number of neurons needed to attain it. Contributions of each cortical area vary for different parameters. Typically, more than 30 randomly sampled neurons were required for an acceptable level of prediction. (G–I) Comparison of the contribution of single units (blue) and multiunits (red) to predictions of hand position, velocity, and gripping force. Single units and multiunits were taken from all cortical areas. Single units' contribution exceeded that of multiunits by approximately 20%. DOI: 10.1371/journal.pbio.0000042.g002

information about any given motor parameter, for each area, different numbers of neurons were required to achieve the same level of prediction. For example, the sample of M1 neurons (33–56 cells) taken alone (Figure 2D–2F) was the best predictor for all motor variables (73% of the variance for hand position, 66% for velocity, 83% for gripping force). The sample of SMA neurons (16–19 cells) produced high predictions of hand position (51%) and velocity (51%), but poor prediction of gripping force (19%). The activity of PMd (64 cells) or S1 (21–39 cells) ensembles predicted hand position (48% for both PMd and S1) and velocity (46% for PMd and 35% for S1) reasonably well, but yielded lower predictions of GF (29% for PMd neurons and 15% for S1). Meanwhile, the sample from PP (63–64 cells) yielded very accurate predictions of gripping force (73%) and hand velocity (52%), but not hand position (25%). Ensemble predictions of gripping force in most cases were more accurate than those obtained from the same population for hand position and velocity. In addition, the ND analysis revealed that predictions of any motor parameter based on combined neural ensemble activity were far superior to those obtained based only on the mean contribution of single neurons.

Another interesting finding emerged from the comparison of the contribution of single-unit versus multiunit activity to the performance of the linear models. Overall, up to 90 single units and 95 multiunits were simultaneously recorded in monkey 1 and 75 single units and 175 multiunits in monkey 2. The cell population was stable not only during the length of the recording sessions but across sessions. The vast majority of the population remained stable for several weeks and, in some cases, months (Nicoletis et al. 2003).

Figure 2G–2I shows that the linear predictions of hand position, velocity, and grasping force were somewhat better when single units were used (by 17%, 20%, and 17%, respectively). That difference could be compensated by increasing the number of channels. For example, as seen in Figure 2G, around 30 additional multiunits compensate for the difference in prediction of hand position provided by 20 single units. That difference was, however, not critical, as the animals could still maintain high levels of BMI performance in all tasks using multiunit activity only. Thus, in contrast to previous studies (Serruya et al. 2002; Taylor et al. 2002) that dealt with fewer motor parameters and a simpler task, we observed that large neuronal samples were needed to achieve a high level of real-time prediction of one or more motor parameters and, consequently, high behavioral proficiency in operating the BMic.

Functional Cortical Reorganization during Operation of BMic

The achievement of high proficiency in the operation of the BMic by the monkeys was consistent with procedural motor learning. Since cortical ensemble recordings were

obtained during behavioral training in all three tasks, it was possible to examine the neurophysiological correlates of this learning process. Both short (within a recording session) and long-term (across recording sessions) physiological modifications took place.

Long-term functional changes in multiple cortical areas were evident in both animals. For instance, the average contribution of single neurons to model performance increased with learning. Figure 3A shows changes in the contribution of single cortical neurons (measured in terms of correlation coefficient, R , color-coded, where blue shows low R ; red, high R) from five cortical areas (PMd, M1, S1, SMA, and M1 ipsilateral) to the linear model that predicted hand position in task 1. Data from 42 recording sessions are shown. In these sessions, predictions of hand position (HPx, HPy) were used to control the cursor on the screen. By the end of the training, very accurate predictions of hand position and velocity were obtained (mean $R \pm$ SEM; HPx = 0.75 ± 0.04 , HPy = 0.72 ± 0.04 , HVx = 0.70 ± 0.03 , and HVy = 0.71 ± 0.02). These high values were reached through a significant increase in contribution of individual neurons to the linear model. When the mean contribution of single neurons was plotted as a function of their cortical area location, differences across cortical areas were found (Figure 3B–3E). The change was higher in SMA (Figure 3E; $R = 0.81$, slope = 0.01 , $p < 0.001$) than in PMd (Figure 3B; $R = 0.81$, slope = 1×10^{-3} , $p < 0.001$), S1 (Figure 3D; $R = 0.67$, slope = 4×10^{-3} , $p < 0.001$), and M1 (Figure 3C; $R = 0.50$, slope = 3×10^{-3} , $p < 0.001$). Note that from the beginning of training, M1 neurons (Figure 3C) provided the highest mean contribution. By the end of 42 sessions, however, the mean contribution of neurons located in other cortical areas (e.g., SMA, PMd, and S1) was as high as that of M1. It is noteworthy that the significant enhancement in contribution occurred for the model predicting hand position (average of all cortical areas, $R = 0.80$, slope = 4×10^{-3} , $p < 0.001$), but not the one predicting hand velocity ($R = 0.05$, slope = 2.2×10^{-4}). This selectivity coincided with the use of a position model in the BMic during these 42 sessions. Thus, long-term training with the BMic using a particular model could result in selective enhancement of the mean contribution of neurons to that model, but not the others.

Changes in Neuronal Direction Tuning during Operation of a BMic

As animals learned to operate the BMic, we also observed short-term changes in neuronal directional tuning, within a single recording session, after switching the BMic mode of operation from pole to brain control. Directional tuning curves (DTCs) reflected dependency of the neuronal firing rate on movement direction of either the pole or the cursor. DTCs were normalized by dividing average firing rates for particular movement directions by the standard deviation of the whole firing rate record and then subtracting the DTC mean. Using that normalization, changes in firing rate with

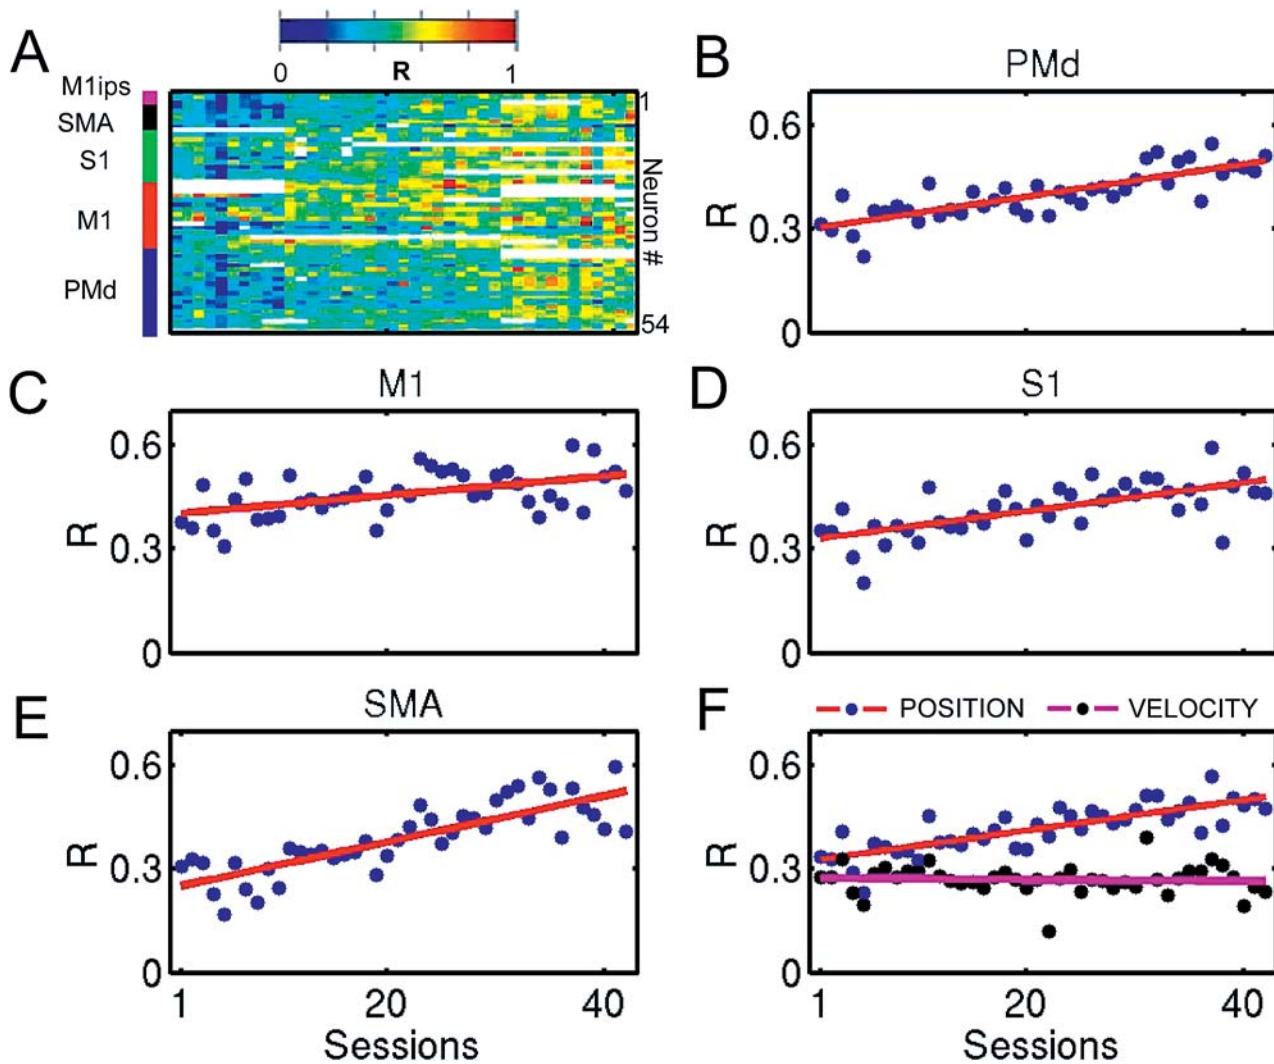


Figure 3. Long-Term Functional Changes in Multiple Cortical Areas

(A) Color-coded (red shows high values; blue, low values) representation of individual contributions measured as the correlation coefficient (R) of neurons to linear model predictions of hand position for 42 training sessions. The average contribution steadily increased with training. The bar on the left indicates cortical location of the neurons.

(B–E) Average contribution of neurons located in different cortical areas (PMd, M1, S1, and SMA, respectively) to hand position prediction during 42 recording sessions.

(F) Average contribution for the whole ensemble to hand position prediction versus hand velocity predictions. A linear increase in contribution was observed only for predictions of hand position.

DOI: 10.1371/journal.pbio.0000042.g003

movement direction were compared with the overall variation of firing rate. Average directional tuning of neural ensembles (DTE) was also assessed, and the spread of the preferred tuning directions was evaluated as the ensemble average of the angles between preferred directions in pairs of neurons. Color-coded population histograms (Figure 4A–4D) displayed the DTCs of all recorded neurons. Polar plots (magenta figures in Figure 4A–4D) showed the DTE and preferred direction spread. Figure 4A–4D shows that DTCs and DTEs for the same neural ensemble had distinct features during pole control (Figure 4A), during brain control with the presence of arm movements (Figure 4B and 4D), and during brain control without arm movements (Figure 4C). Even if the animal was still making arm movements after switching to brain control and direction tuning was calculated in relation to pole movements (compare Figure 4A with 4D), DTC and

DTE changed significantly when compared to curves obtained during pole control ($R = 0.57$ using pole movements as a reference direction, $R = 0.70$ using cursor movements as a reference). The changes in DTC and DTE became greater as the animal ceased to produce arm movements in brain control (Figure 4C) ($R = 0.48$). Notice, however, that the pattern for brain control without arm movements (Figure 4C) was also distinct from that for brain control with arm movements (Figure 4B) ($R = 0.57$). These findings suggest that both the cursor and pole movements influenced the definition of directional tuning profiles in multiple cortical areas.

After the mode of operation was switched to brain control, pole and cursor movements became dissociated. Further, as animals started controlling the BMIc without producing overt hand movements, directional tuning primarily reflected

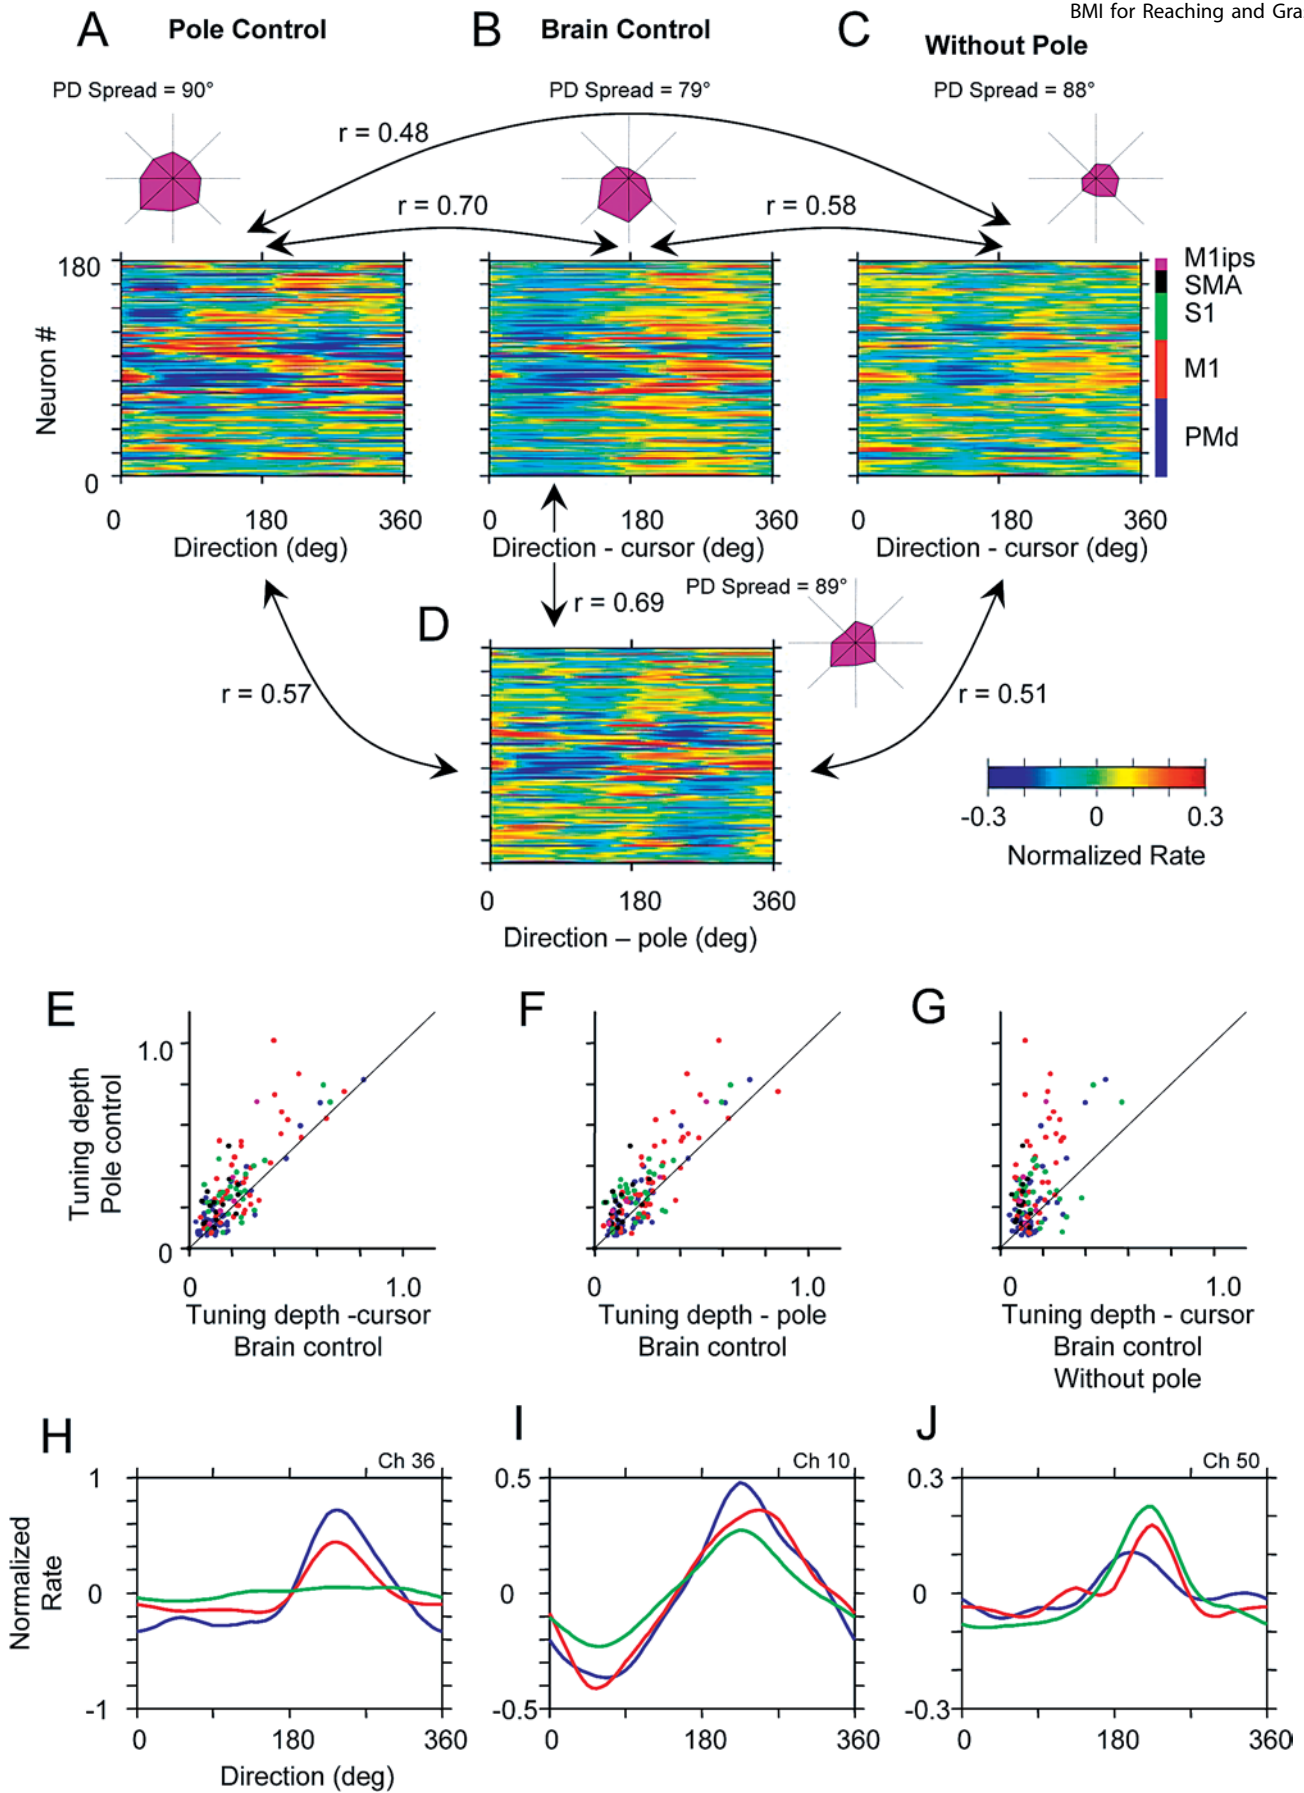


Figure 4. Directional Tuning in Frontoparietal Ensemble during Different Modes of Operation in Task 1

(A–D) Directional tuning during pole control (A), brain control with arm movements (tuning assessed from cursor movements) (B), brain control without arm movements (tuning assessed from cursor movements) (C), and brain control with arm movements (tuning assessed from pole movements) (D). In each of the color-coded diagrams (red shows high values and blue low values; see color scale), the rows depict normalized directional tuning for individual cells. Because of the high directional tuning values of some cells (e.g., that shown in [H]), a color scale limit was set at 0.3 to allow color representation of the largest possible number of cells. Each tuning curve contains eight points that have been interpolated for visual clarity. Correspondence of tuning patterns under different conditions has been quantified using correlation coefficients (shown near lines connecting the diagrams). The highest correspondence was between tuning during pole control and brain control with arm movements. A much less similar pattern of direction tuning emerged during brain control without arm movements. Polar plots near each diagram show average directional tuning for the whole neural ensemble recorded. They indicate an average decrease in tuning strength and shifts in the preferred direction of tuning as the operation mode was switched from pole to brain control. Spread of preferred directions (90° corresponds to uniformly random distribution) is indicated near each polar plot.

(E–G) Scatterplots comparing DTD (maximum minus minimum values of tuning curves) during pole control versus brain control with and without arm movements. DTD during brain control was consistently lower than during pole control. This effect was particularly evident during brain control without arm movements.

(H–J) Changes in directional tuning for individual neurons under different conditions. Blue shows pole control; red, brain control with arm movements (tuning assessed from pole movements); and green, brain control without arm movements. The first illustrated cell (H) was tuned only when the monkey moved its arm, more so during pole control. The second cell (I) had similar tuning during all modes of operation, but tuning depth was the highest for pole control and the lowest for brain control without arm movements. The third cell (J) was better tuned during brain control.

DOI: 10.1371/journal.pbio.0000042.g004

cursor movements. Interestingly, during the transition from pole to brain control, directional tuning depth (DTD) was reduced for most cells. Figure 4E–4G shows this effect by comparing the DTD in individual neurons during pole control (Y axis of Figure 4E–4G) and brain control (X axis). Notice that the reduction in tuning depth was more pronounced when no arm movements were produced during brain control (Figure 4G). Reduction in directional tuning during brain control with no movements characterized 68% of the sampled neurons and included neurons from all cortical areas (see color dots in Figure 4E–4G and examples of Figure 4H and 4I). A small percentage of neurons (14%) did not show such change. Perhaps more surprisingly, a fraction of neurons (18%) enhanced their directional tuning during the switch from pole to brain control (see Figure 4J). These neurons correspond to the dots that are located to the right of the main diagonal in Figure 4E–4G.

Operating the BMIc without making movements was characterized by an appearance of peculiar patterns of directional tuning at the population level. Figure 5A and 5C displays the evolution of DTC and DTE for the same neural ensemble during four task 1 sessions with the robot in the loop. Whereas in each case DTCs during brain control resembled those in pole control, they evolved toward a more organized distribution. Although certain diversity in DTCs remained, clear groups of neurons sharing similar DTCs appeared as a result of training (Figure 5C). Quantitatively, this effect was manifested by a decrease in the spread of preferred directions. This effect was also evident in the polar plots showing population-average tuning (i.e., DTE). The DTE became progressively sharper and rotated clockwise. Throughout the four sessions depicted in Figure 5, tuning depth remained higher during pole than brain control operation of the BMIc (Figure 5B).

Further analysis revealed that significant changes in directional tuning also occurred within a single recording session during brain control (Figure 5D). The session illustrated in Figure 5D was characterized by a gradual improvement in behavioral performance during brain control without arm movements, as evident from measurements made every 5 min (Figure 5E). The population histograms of Figure 5D show that the distribution of DTCs, although variable, became on average tighter across all cortical areas, defining a vertical band across the population histogram. This tightening was

manifested by a decrease in the spread of preferred directions (Figure 5F). Moreover, average tuning depth gradually increased (Figure 5G), but remained lower than that observed during pole control.

Similarity of DTCs during brain control indicated that particular movement directions were associated with simultaneous increases in activity in many neurons; i.e., firing rates of these neurons became more correlated. Indeed, we found increases in broad correlation (100 ms time window) in neuronal firing within and between cortical areas. Thus, during the transition from pole to brain control in task 1, the average correlated firing between pairs of cortical neurons, measured in terms of correlation coefficient (mean \pm SEM), increased from $0.02 \pm 1 \times 10^{-4}$ to $0.06 \pm 2 \times 10^{-4}$, a 3-fold rise that was highly significant (Wilcoxon signed rank test, $p < 0.0001$). All cortical areas tested (M1, PMD, SMA, S1, and M1 ipsilateral) showed increases in correlated firing. The highest within-area increases from pole to brain control were observed in M1 ($\Delta_{\text{Brain-Pole}} = 0.07$), S1 ($\Delta_{\text{Brain-Pole}} = 0.05$), and PMd ($\Delta_{\text{Brain-Pole}} = 0.03$). The highest between-area increases were observed between M1–S1 ($\Delta_{\text{Brain-Pole}} = 0.06$), M1–PMD ($\Delta_{\text{Brain-Pole}} = 0.04$), PMd–S1 ($\Delta_{\text{Brain-Pole}} = 0.04$), M1–SMA ($\Delta_{\text{Brain-Pole}} = 0.02$), and M1_{contra}–M1_{ipsi} ($\Delta_{\text{Brain-Pole}} = 0.02$).

Changes in average firing rates of the neurons during switching from pole to brain control were insubstantial. Firing rates of individual cells ranged from 0.1 to 40 spikes/s (8 ± 8 spikes/s; mean \pm SD). After the mode was switched to brain control and the monkey continued to move the arm, firing rates increased on average by 4% from pole-control level. When the monkey controlled the BMI without moving the arm, the average neuronal firing rates decreased 2.5% from pole-control level.

Real-Time Prediction of Gripping Force

In addition to reproducing hand trajectories with great accuracy, linear models also allowed the reconstruction of fine variations in gripping force produced by both monkeys in tasks 2 and 3. Figure 6A shows that during execution of task 2, most of the recorded cortical neurons contained information about gripping force. In this figure, normalization was achieved by dividing the firing rate of each individual neuron by its standard deviation. In this way, force-related modulations are expressed relative to the

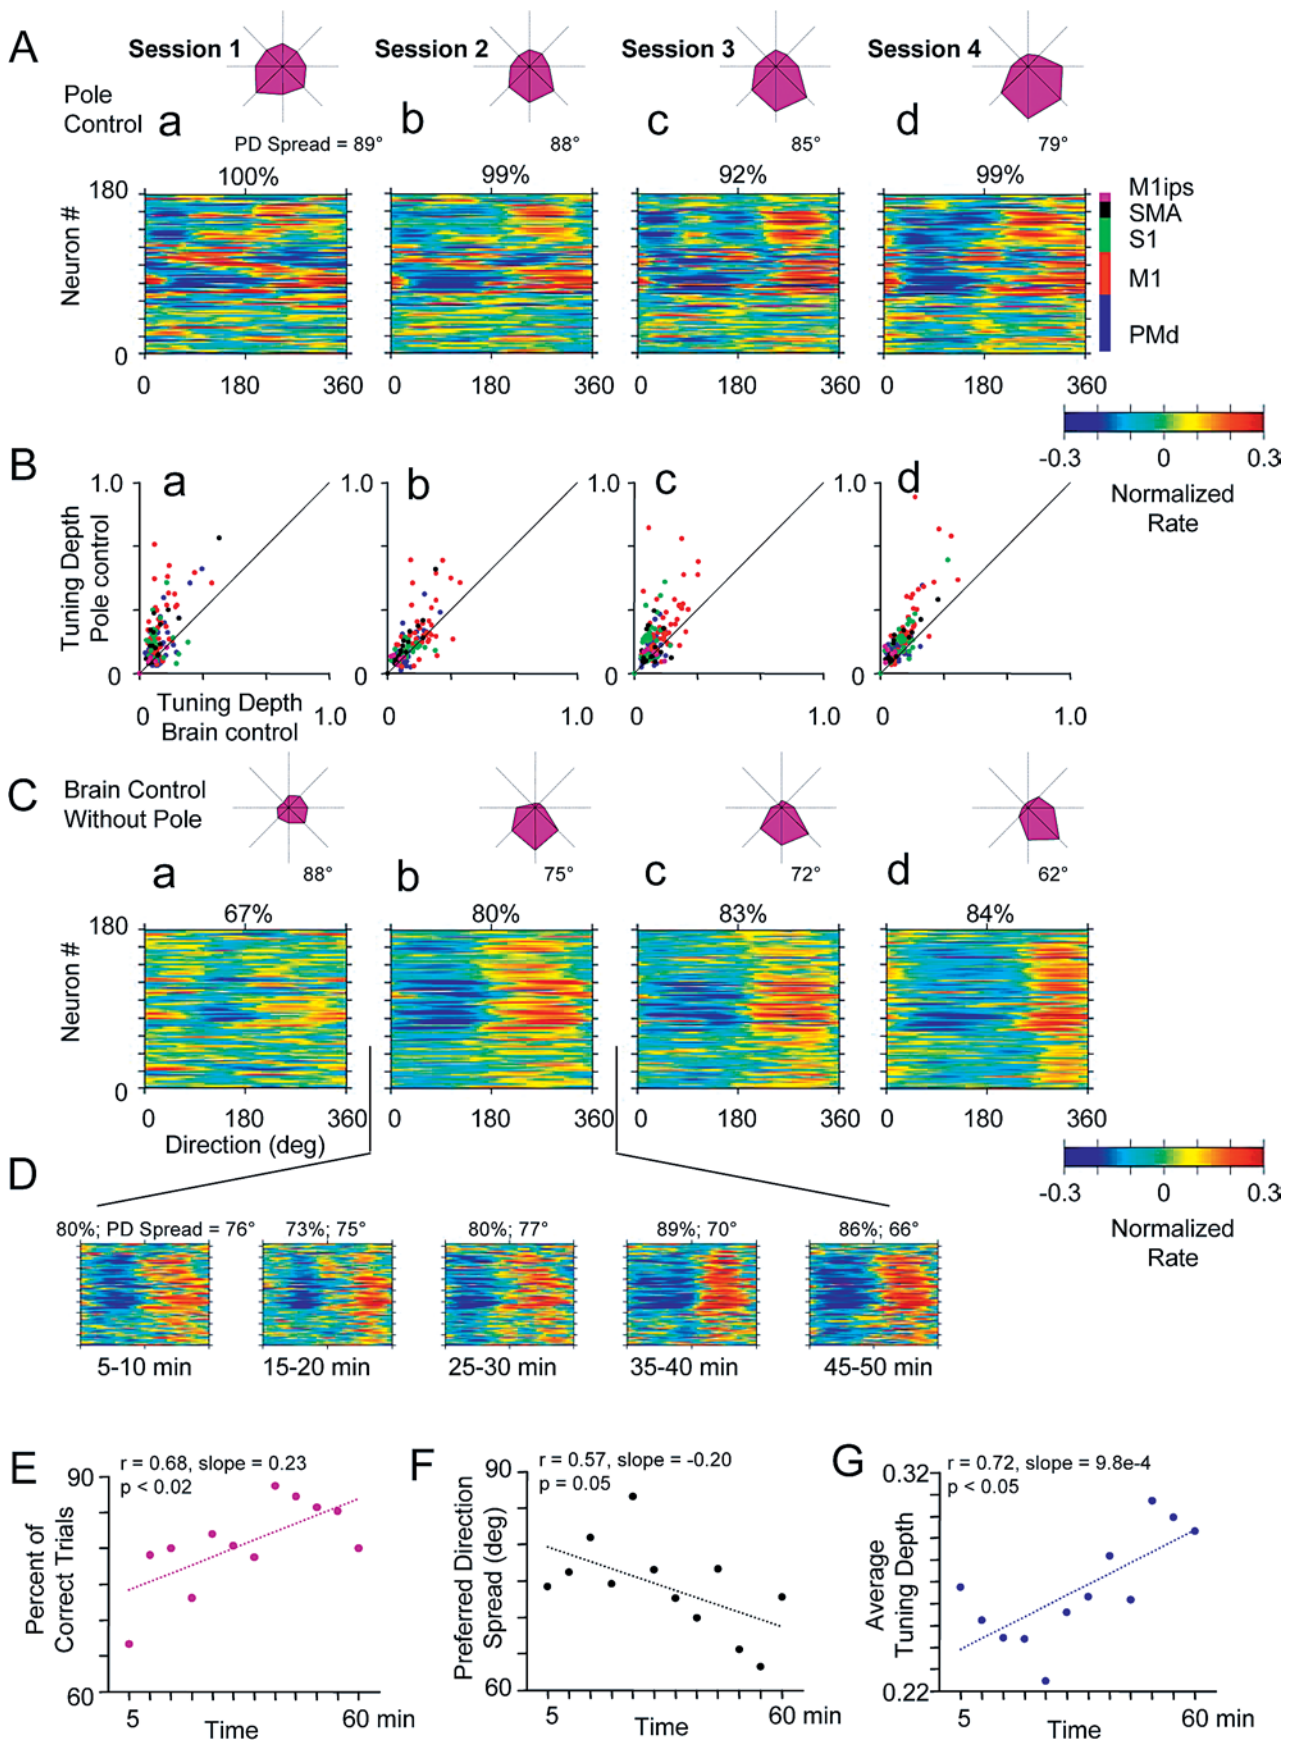


Figure 5. Plasticity of Directional Tuning during Training in Brain Control without Arm Movements

Conventions are as in Figure 4.

(A) Directional tuning profiles during four sessions in pole control (task 1). Percentages of correctly performed trials are indicated for each session.

(B) Scatterplots comparing directional tuning during pole versus brain control for the same sessions. For each day, DTD was on average higher in pole control.

(C) Directional tuning during brain control for the same sessions as in (A). Note the emergence of a population pattern in which a group of neurons (with some exceptions) exhibits a similar preferred direction. This is manifested by a decrease in the spread of preferred directions (shown near polar plots). Notice also a gradual rotation of the population preferred direction (see polar plots) with training.

(D) Gradual changes in DTE during one representative session of brain control without arm movements. This 60-min session was split into 5-min periods, five of which are shown.

(E) Improvement in behavioral performance during a single session (same as in [D]).

(F) Decrease in the spread of preferred directions during that session.

(G) Increase in average tuning depth during the same session.

DOI: 10.1371/journal.pbio.0000042.g005

overall variability of the neuron's firing rate. Both monkeys mastered task 2 in seven to eight sessions. Figure 6B displays the evolution of the average contribution of neurons from different areas of monkey 1 to model predictions during this period. Contribution of contralateral M1 ($R = 0.77$, slope = 0.02, $p < 0.05$) and S1 ($R = 0.85$, slope = 0.02, $p < 0.002$) increased significantly, while that of PMd ($R = 0.19$, slope = 2×10^{-3}), SMA ($R = 0.34$, slope = 0.01), and ipsilateral M1 ($R = 0.38$, slope = -0.01) did not change substantially. For the whole ensemble combined, there was a significant increase in contribution in both monkey 1 ($R = 0.95$, slope = 0.02, $p < 0.001$) and monkey 2 ($R = 0.54$, slope = 0.01, $p < 0.05$). By comparing Figure 3B–3F and Figure 6B, we can see that while M1 and S1 neurons showed changes during both tasks 1 and 2, PMd and SMA neurons showed changes in task 1, but not in task 2. This may reflect the greater involvement of these cortical areas in learning visuomotor spatial relationships than in the production of muscle force.

Switching from pole to brain control did not affect neuronal firing rate correlations in task 2. This could be related to saturation of this parameter because the average firing rate correlation observed during pole control when monkeys performed task 2 (0.056) was already higher than that observed during task 1 (0.022, $p < 0.001$, Wilcoxon rank sum test).

Using a BMic to Reach and Grasp Virtual Objects

Our experiments demonstrated, to our knowledge for the first time, that monkeys can learn to control a BMic to produce a combination of reaching and grasping movements to locate and grasp virtual objects. The major challenge in task 3 was to simultaneously predict hand position and gripping force using the activity recorded from the same neuronal ensemble. This problem could not be reduced to predicting only hand position as in task 1 or gripping force in task 2, because the animal had to sequentially reach and grasp the target.

The DTD of cortical neurons, measured during pole control, increased almost linearly during the learning of task 3 (Figure 6C). Although this effect was significant in all cortical areas tested, its magnitude varied across areas. The most prominent increase in DTD was observed in M1 (red dots in Figure 6C; $R = 0.81$, slope = 0.02). Neurons in S1 (green dots in Figure 6C; $R = 0.82$, slope = 0.02) and ipsilateral M1 (magenta dots in Figure 6C; $R = 0.81$, slope = 0.02) exhibited the next largest increase. Relatively smaller DTD increases were observed in PP (cyan dots in Figure 6C; $R = 0.73$, slope = 0.01), SMA (black dots in Figure 6C; $R = 0.63$,

slope = 0.01), and PMd (blue dots in Figure 6C; $R = 0.51$, slope = 0.01). Similar to task 1, tuning depth was higher during pole control than during brain control. As in task 1, the DTC and DTE patterns changed during training (data not shown). Improvements in model performance occurred as well. Figure 6D and 6E show the evolution in accuracy of real-time predictions of hand position, velocity, and gripping force during 14 sessions for both monkeys. During this period, real-time predictions for both hand position and velocity improved with training, while predictions of gripping force remained high and stable in two monkeys (monkey 1, mean \pm SD, $R = 0.86 \pm 0.04$; monkey 2, $R = 0.79 \pm 0.03$).

The monkeys' performance in brain control in task 3 approximated that during pole control, with characteristic robot displacement (reach) followed by force increase (grasp). Figure 6F and 6G shows several representative examples of reaching and grasping during pole and brain control in task 3 by monkey 1. Hand position (X, Y) and gripping force (F) records are shown. In the display of hand trajectories, the size of the disc at the end of each hand movement shows the gripping force produced by the monkey (Figure 6F) or by the BMic (Figure 6G) to grasp a virtual object. The reach (r) and grasp (g) phases are clearly separated, demonstrating that the monkeys could use the same sample of neurons to produce distinct motor outputs at different moments in time. Thus, during the reaching phase, X and Y changed, while F remained relatively stable. However, as the monkey got closer to the virtual object, F started to increase, while X and Y stabilized to maintain the cursor over the virtual object. Thus, our goal to train the monkey to reproduce coupled sequences of reach-and-grasp movements using the BMic was accomplished.

Discussion

Reliable, long-term operation of a BMic was achieved by extracting multiple motor parameters (i.e., hand position, hand velocity, and gripping force) from the simultaneously recorded activity of frontoparietal neural ensembles. Macaque monkeys learned to use the BMic to reach and grasp virtual objects with a robot even in the absence of overt arm movements. Accurate performance was possible because large populations of neurons from multiple cortical areas were sampled. Thus, the present study shows that large ensembles are preferable for efficient operation of a BMI. This conclusion is consistent with the notion that motor programming and execution are represented in a highly distributed fashion across frontal and parietal areas and that

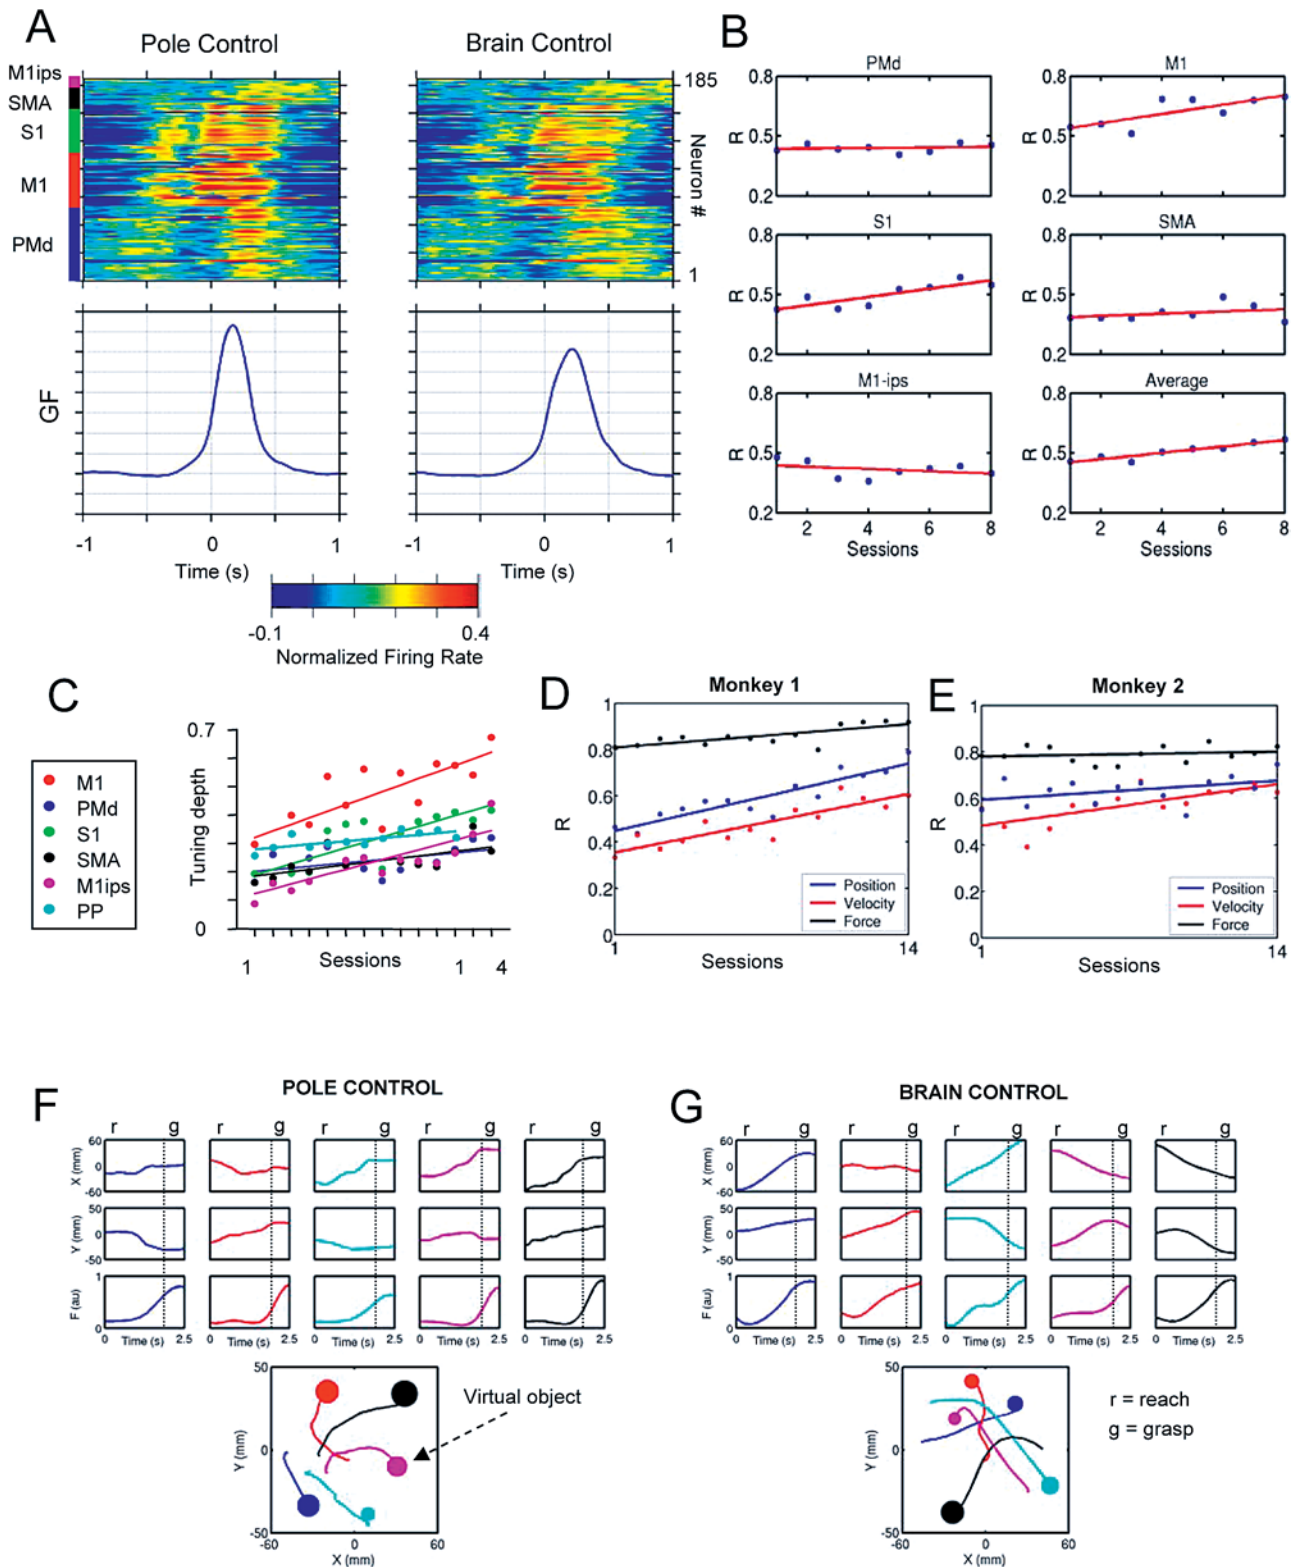


Figure 6. Ensemble Encoding of Gripping Force, Plasticity of Directional Tuning, and Neuronal Contribution to Model Performance during Learning to Control the BMIc for Reaching and Grasping

(A) Perievent time histograms (PETHs) in task 2 for the neuronal population sampled in monkey 1. The plots on top are color-coded (red shows high values; blue, low values). Each horizontal row represents a PETH for a single-neuron or multiunit activity. PETHs have been normalized by subtracting the mean and then dividing by the standard deviation. PETHs are aligned on the gripping force onset (crossing a threshold). Plots at the bottom show the corresponding average traces of gripping force. Note the general similarity of PETHs in pole (left) and brain (right) control in this relatively easy task. Cortical location of neurons is indicated by the bar on the top left. Note the distinct pattern of activation for different areas.

(B) Changes in the mean contribution of neurons from different cortical areas to model predictions during training of monkey 1 in task 2.
 (C) Increases in directional tuning for six cortical areas during training in pole control in task 3.
 (D and E) Increases in neuronal contribution to linear models predicting hand position (blue), hand velocity (red), and gripping force (black) during learning task 3 in both monkeys.
 (F and G) Representative robot trajectories and gripping force profiles in an advanced stage of training in task 3 during both pole and brain control. The bottom graphs show trajectories and the amount of the gripping force developed during grasping each virtual object. The dotted vertical lines in the panels indicate the end of reach phase and the beginning of grasp phase. Note that during both modes of BMIc operation, the patterns of reaching and grasping movements (displacement followed by force increase) were preserved.
 DOI: 10.1371/journal.pbio.0000042.g006

each of these areas contains neurons that represent multiple motor parameters. We suggest that, in principle, any of these areas could be used to operate a BMI, provided that a large enough neuronal sample was obtained. While analysis of ND curves (see Figure 2D–2F) indicates that a significant sample of M1 neurons consistently provides the best predictions of all motor parameters analyzed, neurons in areas such as SMA, S1, PMd, and PP contribute to BMI performance as well.

Our argument for using large neuronal samples is also supported by the fact that some neurons can be lost during chronic recordings, either due to electrode malfunction, modification of electrode tip properties, or simple cell death. Then, a BMI that relies on only very small samples of neurons (e.g., 8–30 cells) (Serruya et al. 2002; Taylor et al. 2002), all derived from a single cortical area, would not be able to provide a broad variety of motor outputs, handle changes in cortical properties, or cope with alterations in the neuronal sample over time.

Another important finding of this study is that accurate real-time prediction of all motor parameters as well as a high level of BMI control can be obtained from multiunit signals. This observation is essential because it eliminates the need to develop elaborate real-time spike-sorting algorithms, a major technological challenge, for the design of a future cortical neuroprosthesis for clinical applications.

Our experiments also demonstrate that the initial introduction of a mechanical device, such as the robot arm, in the control loop of a BMIc significantly impacts learning and task performance. After the robot was introduced in the control loop, the monkey had to adjust to the dynamics of this artificial actuator. As a result, there was an immediate drop in performance (see Figure 3G). With further training, however, the animals were able to overcome the difficulties. Thus, the simple utilization of the output of a real-time model to move a cursor on a computer screen (Serruya et al. 2002; Taylor et al. 2002) does not fully test the limitations and challenges involved in operating a clinically relevant BMI. Instead, such testing must include the incorporation in the apparatus of the mechanical actuator designed to enact the subject's motor intentions and training the subject to operate it.

As was proposed recently (Nicoletis and Ribeiro 2002), multisite, multielectrode recordings (Nicoletis et al. 2003) also allowed us to quantify neurophysiological modifications occurring in cortical networks, as animals learned motor tasks of different complexity. At a single-neuron level, one modification observed was a reduction in the strength of directional tuning as animals switched from pole to brain control of the BMI, an effect that reached its maximum as animals ceased to produce overt arm movements. This finding touches directly on the ongoing debate of two opposing views of what the motor cortex encodes (Mussa-Ivaldi 1988; Georgopoulos et al. 1989; Kakei et al. 1999; Todorov 2000; Johnson et al. 2001; Scott et al. 2001). At the

first glance, the reduction in tuning depth in the absence of arm movements could be interpreted as providing support to the notion that directional tuning in the motor cortex is highly influenced by movement dynamics. Thus, the elimination of proprioceptive feedback during brain control could explain the significant reduction in directional tuning. However, a smaller but significant decrease in directional tuning was also observed during brain control while animals still used their hands to move the pole. This suggests that directional tuning reflects neither movement dynamics nor abstract motor goals alone, but rather their combination. Additional findings support this contention. For example, a small fraction of M1 and S1 neurons became better directionally tuned when the monkey did not make hand movements during brain control (see Figure 4Ba–4Bd and Figure 5G). Moreover, during brain control there was a significant increase in pairwise-correlated firing and a tendency for groups of neurons to exhibit rather similar DTCs (see Figures 4 and 5). Increases in tuning depth accompanied improvements in performance during brain control, although values observed during pole control were never reached (see Figure 5E and 5G). All together, these physiological changes suggest that as animals learn to operate the BMI during brain control, visual feedback signals representing the goal of movement, rather than information about arm movements per se, become the main guiding signal to the cortical neurons that drive the BMIc. Thus, we hypothesize that, as monkeys learn to formulate a much more abstract strategy to achieve the goal of moving the cursor to a target, without moving their own arms, the dynamics of the robot arm (reflected by the cursor movements) become incorporated into multiple cortical representations. In other words, we propose that the gradual increase in behavioral performance during brain control of the BMI emerged as a consequence of a plastic reorganization whose main outcome was the assimilation of the dynamics of an artificial actuator into the physiological properties of frontoparietal neurons. This hypothesis is consistent with previous observations in paralyzed humans who learned to move a cursor on the screen using cortical activity (Kennedy and King 2000). It is also supported by the results that cortical neurons may modulate their firing rate either during use of tools (Iriki et al. 1996), according to cursor movement on the screen rather than underlying arm movements (Alexander and Crutcher 1990; Shen and Alexander 1997), or in relation to the orientation of spatial attention (Lebedev and Wise 2001).

Our results on cortical reorganization are very distinct from a previous claim of plastic changes in directional tuning of cortical cells during the use of BMI (Taylor et al. 2002). First, in that previous report, the population vector model yielded poor predictions when applied to activity of a small sample ($n = 18$) of M1 cells. Introduction of visual feedback improved the subject's performance to a point in which

monkeys could use a BMI to produce stereotypic center-out movements of a cursor. The authors claimed that changes in cell-preferred direction occurred after switching to brain control. However, preferred directions were derived not from the real-movement directions of the hand or the cursor, but rather from ideal directions defined by target locations. In addition, a wide 420–990 ms time window was used to measure firing rates. This window was comparable to movement duration. Therefore, differences in movement trajectories and duration between hand and brain control, rather than true differences in cell directional tuning, could lead to different estimates of preferred direction. The report also claims that tuning strength increased with training in brain, but not hand, control. However, tuning depth was evaluated by measuring covariation between firing rate modulations and target locations, rather than actual movement trajectories. Because during training, cursor trajectories gradually approached a straight line connecting the starting point and the target, the observed improvement in covariation between target locations and neuronal firing rate modulations could simply reflect the improvement in movement accuracy. These considerations should be taken into account to decide how much of the plasticity reported by Taylor et al. (2002) reflects real cortical reorganization instead of resulting from the improvement in the animal's behavioral performance during the task used to measure directional tuning.

In the present study, all the changes in firing and tuning properties of neuronal ensembles were related to the actual trajectories produced by the monkeys during pole and brain control. Moreover, the relationship between the neuronal firing and movement patterns was evaluated continuously. Thus, the cortical changes reported here more closely reflected the relationship between neuronal signals and motor behaviors that they underlie.

Overall, the present findings demonstrate that it is reasonable to envision that a cortical neuroprosthesis for restoring upper-limb movements could be implemented in the future, following the basic BMIc principles described here. We propose that long-term operation of such a device by paralyzed subjects would lead, through a process of cortical plasticity, to the incorporation of artificial actuator dynamics into multiple brain representations. Ultimately, we predict that this assimilation process will not only ensure proficient operation of the neuroprosthesis, but it will also confer to subjects the perception that such apparatus has become an integral part of their own bodies.

Materials and Methods

Behavioral training and electrophysiology. Two adult female monkeys (*Macaca mulatta*) were used in this study. All procedures conformed to the National Research Council's *Guide for the Care and Use of Laboratory Animals* (1996) and were approved by the Duke University Animal Care and Use Committee.

At the time of surgery, animals had completed a period of chair training, and one of them (monkey 2) was familiarized with the requirements of task 1 (a large target size was used in this preliminary training). Multiple arrays containing 16–64 microwires each (separation between adjacent microwires = 300 μm) were implanted in several frontal and parietal cortical areas in each animal (Nicoletis et al. 2003) (96 in monkey 1 and 320 in monkey 2). Implanted areas included the dorsal premotor cortex (PMd), supplementary motor area (SMA), and the primary motor cortex (M1) in both hemispheres. In monkey 1 an implant was placed in the primary somatosensory cortex (S1). In monkey 2, the medial intraparietal area (MIP) of the

posterior parietal cortex (PP) was also implanted. The monkeys performed the tasks with their left arms, which were contralateral to the areas with the best cell yield. Upon recovery from this surgical procedure, animals were transferred to the experimental apparatus illustrated in Figure 1A and started behavioral training.

Monkeys were seated in a primate chair facing a computer monitor. They were trained to perform three different tasks using a hand-held pole equipped with a pressure transducer (PCB Piezo-electrics Inc., Depew, New York, United States) for measuring grasping force. The position of the monkey's hand was obtained from an infrared marker located on top of the pole. The marker was monitored by an optical tracking system (Optotrak, Northern Digital, Waterloo, Ontario, Canada). In the first task, the monkeys were shown a small disk (the "cursor") and a larger disk (the "target"). They had to use the pole to put the cursor over the target and remain within it for 150 ms. Should the monkey cross the target too fast, the target disappeared and the trial was not rewarded. Each trial began with a target presented in a random position on the screen. The monkeys had 5 s to hit the target and get a juice reward. In the second task, the monkeys were presented with the cursor in the center of the screen and two concentric circles. The ring formed by these two circles instructed the amount of gripping force the animals had to produce. The pole was fixed, and the cursor grew in size as the monkey gripped the pole, providing continuous visual feedback of the gripping force. Force instruction changed every trial. The third task contained elements of tasks 1 and 2. In this task, the monkeys were presented with the cursor, the target, and the force-instructing ring and were required to manipulate the pole to put the cursor over the target and match the ring size by developing the proper amount of gripping force, while staying inside the target. The monkeys received juice rewards for correct performance. In task 1, the monkeys were initially trained without the robot in the loop, but after a certain number of sessions, the robot was incorporated to the loop. In tasks 2 and 3, the robot was always part of the loop.

A 512-multichannel acquisition processor (Plexon Inc., Dallas, Texas, United States) was employed to simultaneously record from single neuron and multiunit activity during each recording session. EMG signals were recorded from the skin surface just above the belly of the wrist flexors, wrist extensors, and biceps muscles using gold disc electrodes (Grass Instrument Co., West Warwick, Rhode Island, United States) filled with conductive cream. These signals were amplified (gain, 10,000 \times), high-pass filtered, rectified, and smoothed (kernel convolution).

Linear model. Hand position, velocity, and gripping force were modeled as a weighted linear combination of neuronal activity using a multidimensional linear regression or Wiener filter, the basic form of which is

$$\mathbf{y}(t) = \mathbf{b} + \sum_{u=-m}^n \mathbf{a}(u)\mathbf{x}(t-u) + \boldsymbol{\varepsilon}(t)$$

In this equation, $\mathbf{x}(t-u)$ is an input vector of neuronal firing rates at time t and timelag u , $\mathbf{y}(t)$ is a vector of kinematic and dynamic variables (e.g., position, velocity, gripping force) at time t , $\mathbf{a}(u)$ is a vector of weights at timelag u , \mathbf{b} is a vector of y-intercepts, and $\boldsymbol{\varepsilon}(t)$ are the residual errors. The lags in the summation can in general be negative (in the past) or positive (in the future) with respect to the present time t . We only considered lags into the past.

This equation can be recast in matrix form as

$$\mathbf{Y} = \mathbf{X}\mathbf{A},$$

here each row in each matrix is a unit of time and each column is a data vector. Note that matrix \mathbf{X} contains lagged data and thus has a column for each lag multiplied by the number of channels; e.g., 100 channels and 10 lags imply 1000 columns. The y-intercept is handled by prepending a column of ones to matrix \mathbf{X} . Matrix \mathbf{A} is then solved by

$$\mathbf{A} = \text{inv}(\mathbf{X}^T\mathbf{X})\mathbf{X}^T\mathbf{Y}$$

Real-time predictions of motor parameters. Predictions of hand trajectory and grasping force were generated using the Wiener filter described above. Neuronal firing rates were sampled using 100 ms bins, and 10 bins preceding a given point in time were used for training the model and predicting with it. Models were trained with 10 min of data and tested by applying them to subsequent records. In individual neuron analysis, a model was trained using single-unit/multiunit activity only and then tested for predictions of motor parameters. In velocity mode, the model was trained using velocity

derived from position measurements by digital differentiation. During brain control, predicted velocity was digitally integrated to provide an output position signal. To avoid slow drifts, this signal was high-pass filtered with a first-order Butterworth filter (cutoff frequency of 2 Hz). The linear models independently predicted X, Y, and Z hand position coordinates. However, because the three tasks reported in this study took place in the X–Y plane, the predictions of position and velocity along the Z axis were not used. Several alternative decoding algorithms were tested offline, including a Kalman filter, normalized least-mean squares filter, and a feed-forward backpropagation artificial neural network. None of these methods could consistently outperform the Wiener filter.

Robot arm and gripper. A 6 DOF robotic arm equipped with a 1 DOF gripper (The ARM, Exact Dynamics, Didam, The Netherlands) was used in this study. The gripper was sensorized with pressure transducers (Flexiforce, Tekscan, Boston, Massachusetts, United States) of 1 lb (2.2 kg) force range for providing grasping force feedback. Position feedback of the robot was obtained through the built-in joint encoders. Both the commands for controlling the robot and the feedback were in Cartesian coordinates. The communication between the client computer and the robot was performed via the CAN bus (National Instruments, Austin, Texas, United States) (sampling period, 60 ms). For the tasks involving grasping, the gripper had a light object inserted made of foam material. This object was squeezed by the gripper in proportion to either the force applied by the monkey in the pole or to the brain signal.

Data analysis. The monkeys' behavior was continuously monitored and videotaped throughout each recording session. The percentage of correctly performed trials and the time to accomplish each trial, during both pole and brain controls, were used as measures of performance. Chance performance for each task was determined using Monte Carlo simulations of a random walk with 2, 1, and 3 DOF (for task 1, task 2, and task 3, respectively.) The velocity of the random walk was varied from 1 to 500 mm/s, with 10,000 trials for each velocity. For tasks 1 and 3, this was the velocity of the cursor; for task 2, velocity corresponded to the rate in change of the cursor radius. Because each monkey operated the pole at different speed, predictions shown in Figure 1C–1E are based on the average velocity across all sessions for a given monkey.

Random neuron-dropping (ND) technique was implemented as described by Wessberg et al. (2000). Population data (10 min) were used to fit a linear model, which was used to predict motor parameters from the subsequent record. A single neuron was then randomly removed from the population, the model retrained, and new predictions generated. This process was repeated until only one neuron remained. The average squared linear correlation coefficient (R^2) as a function of number of neurons was obtained by repeating this procedure 30 times for each ensemble. Curves were obtained with populations of neurons segregated from M1, PMd, S1, SMA, and PP and for single units and multiunits.

References

- Alexander GE, Crutcher MD (1990) Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. *J Neurophysiol* 64: 164–178.
- Birbaumer N, Ghanayim N, Hinterberger T, Iversen I, Kotchoubey B, et al. (1999) A spelling device for the paralysed. *Nature* 398: 297–298.
- Bomze HM, Bulsara KR, Iskandar BJ, Caroni P, Skene JH (2001) Spinal axon regeneration evoked by replacing two growth cone proteins in adult neurons. *Nat Neurosci* 4: 38–43.
- Bunge MB (2001) Bridging areas of injury in the spinal cord. *Neuroscientist* 7: 325–339.
- Burnod Y, Baraduc P, Battaglia-Mayer A, Guigon E, Koehlin E, et al. (1999) Parieto-frontal coding of reaching: An integrated framework. *Exp Brain Res* 129: 325–346.
- Chapin JK, Moxon KA, Markowitz RS, Nicolelis MAL (1999) Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nat Neurosci* 2: 664–670.
- Donoghue JP (2002) Connecting cortex to machines: Recent advances in brain interfaces. *Nat Neurosci Supp* 5: 1085–1088.
- Fetz EE (1969) Operant conditioning of cortical unit activity. *Science* 163: 955–957.
- Fetz EE (1992) Are movement parameters recognizably coded in activity of single neurons? *Behav Brain Sci* 15: 679–690.
- Fetz EE, Baker MA (1973) Operantly conditioned patterns of precentral unit activity and correlated responses in adjacent cells and contralateral muscles. *J Neurophysiol* 36: 179–204.
- Fetz EE, Finocchio DV (1971) Operant conditioning of specific patterns of neural and muscular activity. *Science* 174: 431–435.
- Fetz EE, Finocchio DV (1975) Correlations between activity of motor cortex

cells and arm muscles during operantly conditioned response patterns. *Exp Brain Res* 23: 217–240.

Georgopoulos AP, Lurito JT, Petrides M, Schwartz AB, Massey JT (1989) Mental rotation of the neuronal population vector. *Science* 243: 234–236.

Johnson MTV, Mason CR, Ebner TJ (2001) Central processes for the multiparametric control of arm movements in primates. *Curr Opin Neurobiol* 11: 684–688.

Kakei S, Hoffman DS, Strick PL (1999) Muscle and movement representations in the primary motor cortex. *Science* 285: 2136–2139.

Kennedy PR, King B (2000) Dynamic interplay of neural signals during the emergence of cursor related cortex in a human implanted with the neurotrophic electrode. In: Chapin J, Moxon K, editors. *Neural prostheses for restoration of sensory and motor functions*. Boca Raton, Florida: CRC Press. pp. 221–233.

Iriki A, Tanaka M, Iwamura Y (1996) Coding of modified body schema during tool use by macaque postcentral neurones. *Neuroreport* 7: 2325–2330.

Lebedev MA, Wise SP (2001) Tuning for the orientation of spatial attention in dorsal premotor cortex. *Eur J Neurosci* 13: 1002–1008.

Li CS, Padoa-Schioppa C, Bizzi E (2001) Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* 30: 593–607.

Messier J, Kalaska JF (2000) Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task. *J Neurophysiol* 84: 152–165.

Mussa-Ivaldi FA (1988) Do neurons in the motor cortex encode movement direction?: An alternative hypothesis. *Neurosci Lett* 91: 106–111.

National Research Council (1996) *Guide for the care and use of laboratory animals*. Rev. ed. Washington, District of Columbia: National Academy Press. 140 p.

Acknowledgments

This work was supported by grants from the Defense Advanced Research Projects Agency and the James S. McDonnell Foundation to MALN and by the National Science Foundation to MALN and CSH. We thank Gary Lehew, Aaron Sandler, Ming Qian, Erin Phelps, Kristen Shanklen, and Susan Halkiotis for their outstanding technical assistance.

Conflicts of Interest. The authors have declared that no conflicts of interest exist.

Author Contributions. JMC and MALN conceived and designed the experiments. JMC, MAL, and MALN performed the experiments. JMC integrated the robotic interface. JMC and MAL analyzed the data. JEO and CSH contributed Beowulf cluster simulations and analyses. REC contributed the initial stages of this study. JMC, DMS, and MAL recorded the EMGs. DFD performed the surgeries; JMC, MAL, REC, PGP, and MALN assisted in the surgeries. JMC, MAL, and MALN wrote the paper.



- Nicolelis MAL (2001) Actions from thoughts. *Nature* 409: 403–407.
- Nicolelis MAL (2003) Brain–machine interfaces to restore motor function and probe neural circuits. *Nat Rev Neurosci* 4: 417–422.
- Nicolelis MAL, Ribeiro S (2002) Multi-electrode recordings: The next steps. *Curr Opin Neurobiol* 12: 602–606.
- Nicolelis MAL, Dimitrov D, Carmena JM, Crist R, Lehew G, et al. (2003) Chronic, multi-site, multi-electrode recordings in macaque monkeys. *Proc Natl Acad Sci U S A* 100: 11041–11046.
- Nobunaga AI, Go BK, Karunas RB (1999) Recent demographic and injury trends in people served by the model spinal cord injury care systems. *Arch Phys Med Rehabil* 80: 1372–1382.
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci* 5: 805–811.
- Ramon-Cueto A, Plant GW, Avila J, Bunge MB (1998) Long-distance axonal regeneration in the transected adult rat spinal cord is promoted by olfactory ensheathing glia transplants. *J Neurosci* 18: 3803–3815.
- Schmidt EM (1980) Single neuron recording from motor cortex as a possible source of signals for control of external devices. *Ann Biom Eng* 8: 339–349.
- Schwab ME (2002) Repairing the injured spinal cord. *Science* 295: 1029–1031.
- Scott SH, Gribble PL, Graham KM, Cabel DW (2001) Dissociation between hand motion and population vectors from neural activity in motor cortex. *Nature* 413: 161–165.
- Serruya MD, Hatsopoulos NG, Paninski L, Fellows MR, Donoghue JP (2002) Instant neural control of a movement signal. *Nature* 416: 141–142.
- Shen L, Alexander GE (1997) Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. *J Neurophysiol* 77: 1195–1212.
- Talwar SK, Xu S, Hawley ES, Weiss SA, Moxon KA, et al. (2002) Rat navigation guided by remote control. *Nature* 417: 37–38.
- Taylor DM, Tillery SI, Schwartz AB (2002) Direct cortical control of 3D neuroprosthetic devices. *Science* 296: 1829–1832.
- Todorov E (2000) Direct cortical control of muscle activation in voluntary arm movements: A model. *Nat Neurosci* 3: 391–398.
- Uchida N, Buck DW, He D, Reitsma MJ, Masek M, et al. (2000) Direct isolation of human central nervous system stem cells. *Proc Natl Acad Sci U S A* 97: 14720–14725.
- Wessberg J, Nicolelis MAL (2003) Optimizing a linear algorithm for real-time robotic control using chronic cortical ensemble in recording monkeys. *J Cogn Neurosci*. In press.
- Wessberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, et al. (2000) Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408: 361–365.
- Wise SP, Boussaoud D, Johnson PB, Caminiti R (1997) Premotor and parietal cortex: Corticocortical connectivity and combinatorial computations. *Annu Rev Neurosci* 20: 25–42.